

EXHIBIT N

INTRODUCTION

Guide to Anthropic's prompt engineering resources

Getting access to Claude

Getting started with Claude

Your first chat with Claude

Configuring GPT prompts for Claude

Claude for Google Sheets

Glossary

PROMPT DESIGN

Introduction to prompt design

Constructing a prompt

Optimizing your prompt

USEFUL HACKS

Let Claude say "I don't know" to prevent hallucinations

Give Claude room to "think" before responding

Ask Claude to think step-by-step

Break complex tasks into subtasks

Prompt Chaining

Check Claude's comprehension

Ask Claude for rewrites

USE CASES

Content Generation

Introduction to prompt design

Claude is trained to be a helpful, honest, and harmless assistant. It is used to speaking in *dialogue* and you can instruct it in regular English.

The quality of the instructions you give Claude can have a large effect on the quality of its outputs, especially for complex tasks. This guide to prompt design will help you learn how to craft prompts that produce accurate and consistent results.

Claude works by sequence prediction

Claude is a conversational assistant, based on a large language model. The model uses all the text that you have sent it (your prompt) and all the text it has generated so far to predict the next token that would be most helpful.

This means that Claude constructs its responses one set of characters at a time, in order. It cannot go back and edit its response after it has written it unless you give it a chance to do so in a subsequent prompt.

Claude can also only see (and make predictions on) what is in its context window. It can't remember previous conversations unless you put them in the prompt, and it can't open links.

What is a prompt?

The text that you give Claude is designed to elicit, or "prompt", a relevant output. A prompt is usually in the form of a question or instructions. For example:

Prompt

Human: Why is the sky blue?

Assistant:

The text that Claude responds with is called a "output".

Output

The sky appears blue to us due to how the molecules in the air scatter light from the sun. The shorter wavelength light (blue) scatters more easily than the longer wavelength light (like reds and yellows), so when we look up at the sky, what we see is all the scattered blue light. This is known as Rayleigh scattering. The color of the sky varies based on conditions like the angle of the sun and amount of dust/water in the air, but during the day with a clear sky, the blue wavelength of light is most responsible for the color we perceive.

Human: / Assistant: formatting

Claude is trained to fill in text for the Assistant role as part of an ongoing dialogue between a human user (Human:) and an AI assistant (Assistant:).

Prompts sent via the API must contain `\n\nHuman:` and `\n\nAssistant:` as the signals of who's speaking. In Slack and our web interface we automatically add these for you.

For example, in the **Console** or in **Claude-in-Slack**, you can just ask Claude:

Prompt in Console or Claude-in-Slack

In one sentence, what is good about the color blue?

And it will respond:

Output

Blue is often seen as a calming and soothing color.

If you send the same prompt to the **API**, it may behave in unexpected ways, like making up answers well beyond what was asked for in the prompt. This is because Claude is trained to fill in text for the Assistant role as part of an ongoing dialogue between a human user (Human:) and an AI assistant (Assistant:). Without this structure, Claude doesn't know what to do or when to stop, so it just keeps on going with the arc that's already present.

The prompt sent to the API must be:

Prompt for API

Human: In one sentence, what is good about the color blue?

Assistant:

TABLE OF CONTENTS

Claude works by sequence prediction

What is a prompt?

Human: / Assistant: formatting

Prompt length

💡 Why?

Claude has been trained and fine-tuned using RLHF (reinforcement learning with human feedback) methods on `\n\nHuman:` and `\n\nAssistant:` data like this, so **you will need to use these prompts in the API** in order to stay "on-distribution" and get the expected results. It's important to remember to have the two newlines before both Human and Assistant, as that's what it was trained on.

Prompt length

The maximum prompt length that Claude can see is its [context window](#). Claude's context window is currently ~75,000 words / ~100,000 tokens / ~340,000 Unicode characters.

🕒 Updated 12 days ago

← [Glossary](#)

[Constructing a prompt](#) →

Did this page help you? ☐ Yes ☐ No

🌐 English